arXiv:1604.00204v1 [cs.CR] 1 Apr 2016

Verifying Security Policies using Host Attributes

Cornelius Diekmann[1], Stephan-A. Posselt[1], Heiko Niedermayer[1], Holger Kinkelin[1], Oliver Hanka[2], and Georg Carle[1]

[1] Technische Universität München surname@net.in.tum.de
[2] Airbus Group Innovations first_name.surname@eads.net

Abstract. For the formal verification of a network security policy, it is crucial to express the verification goals. These formal goals, called security invariants, should be easy to express for the end user. Focusing on access control and information flow security strategies, this work discovers and proves universal insights about security invariants. This enables secure and convenient auto-completion of host attribute configurations. We demonstrate our results in a civil aviation scenario. All results are machine-verified with the Isabelle/HOL theorem prover.

1 Introduction

A distributed system, from a networking point of view, is essentially a set of interconnected hosts. Its connectivity structure comprises an important aspect of its overall attack surface, which can be dramatically decreased by giving each host only the necessary access rights. Hence, it is common to protect networks using firewalls and other forms of enforcing network level access policies. However, raw sets of such policy rules e. g., firewall rules, ACLs, or access control matrices, scale quadratically with the number of hosts and "controlling complexity is a core problem in information security" [15]. A case study, conducted in this paper, reveals that even a policy with only 10 entities may cause difficulties for experienced administrators. Expressive policy languages can help to reduce the complexity. However, the question whether a policy fulfills certain security invariants and how to express these often remains.

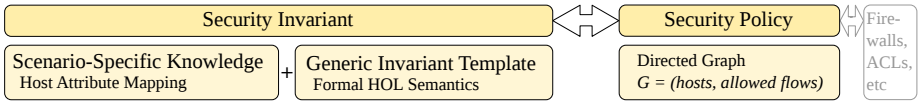

Fig. 1: Formal objects: Security invariant and security policy.

Using an attribute-based [21] approach, we model simple, static, positive security policies with expressive, Mandatory Access Control (MAC) security invariants. The formal objects, illustrated in Fig. 1, are carefully constructed for their use-case. The policy is simply a graph, which can for example be extracted from

or translated to firewall rules. The security invariants are split into the formal semantics, accessible to formal analysis, and scenario-specific knowledge, easily configurable by the end user. This model landscape enables verification of security policies. Primarily, we contribute the following universal insights for constructing security invariants.

1. Both provably *secure* and *permissive* default values for host attributes can be found. This auto completion decreases the user's configuration effort.
2. The security strategy, information flow or access control, determines whether a security violation occurs at the sender's or at the receiver's side.
3. A violated invariant can always be repaired by tightening the policy *if and only if* the invariant holds for the deny-all policy.

We formally introduce the underlying model in Sect. 2. Then we present three examples of security invariant templates in Sect. 3 and conduct a formal analysis in Sect. 4. Our implementation and a case study are presented in Sections 5 and 6. Related work is described in Sect. 7. We conclude in Sect. 8.

2 Formal Model

We formalized all our theory in the Isabelle/HOL theorem prover [20]. To stay focused, we omit all proofs in this document but point the reader to the complete formal proofs by roman reference marks. For example, when the paper states 'foo[iv]', the machine-verified proof for the claim 'foo' can be found by following the corresponding endnote. Note that standards such as Common Criteria [7] require formal verification for their highest *Evaluation Assurance Level* (EAL7) and the Isabelle/HOL theorem prover is suitable for this purpose [7, §A.5]. Therefore, our approach is not only suitable for verification, but also a first step towards certification.

Retaining network terminology, we will use the term *host* for any entity which may appear in a policy*, e. g., collections of IP addresses, names, or even roles. A *security policy* is "a specific statement of what is and is not allowed" [5]. Narrowing its scope to network level access control, a security policy is a set of rules which state the allowed communication relationships between hosts. It can be represented as a directed graph.

Definition 1 (Security Policy). *A security policy is a directed graph* $G = (V, E)$*, where the hosts* V *are a set of type* $\mathcal{V}$ *and the allowed flows* E *are a set of type* $\mathcal{V}\times\mathcal{V}$*. The type of* G *is abbreviated by* $\mathcal{G} = (\mathcal{V}\ set) \times ((\mathcal{V}\times\mathcal{V})\ set)$*.*

A policy defines rules (*"how?"*). It does not justify the intention behind these rules (*"why?"*). To reflect the *why?*-question, we note that depending on a concrete scenario, hosts may have varying security-relevant attributes. We model a host attribute of arbitrary type Ψ and establish a total mapping from the hosts

*In contrast to common policy terminology, we do not differentiate between subjects and targets (objects) as they are usually indistinguishable on the network layer and a host may act as both.

V to their scenario-specific attribute. Security invariants can be constructed by combining a *host mapping* with a *security invariant template*. Latter two are defined together because the same Ψ is needed for a related host mapping and security invariant template. Different Ψ may appear across several security invariants.

Definition 2 (Host Mapping and Security Invariant Template). *For scenario-specific attributes of type Ψ, a host mapping P is a total function which maps a host to an attribute. P is of type $\mathcal{V} \Rightarrow \Psi$.*

A security invariant template m is a predicate $m(\mathcal{G}, (\mathcal{V} \Rightarrow \Psi))$, defining the formal semantics of a security invariant. Its first argument is a security policy, its second argument a host attribute mapping. The predicate $m(G, P)$ returns true iff the security policy G fulfills the security invariant specified by m and P.*

Example 1. Label-based information flow security can be modeled with a simplified version of the Bell LaPadula model [2,3]. Labels, more precisely *security clearances*, are host attributes $\Psi = \{unclassified, confidential, secret, topsecret\}$. The Bell LaPadula's no read-up and no write-down rules can be summarized by requiring that the security clearance of a receiver r should be greater-equal than the security clearance of the sender s, for all $(s, r) \in E$. With a total order '$\leq$' on Ψ, the security invariant template can be defined as $m((V, E), P) \equiv \forall (s, r) \in E.\ P(s) \leq P(r)$.

Let the scenario-specific knowledge be that database $db_1 \in V$ is *confidential* and all other hosts are *unclassified*. Using lambda calculus, the total function P can be defined as ($\lambda h.$ **if** $h = db_1$ **then** *confidential* **else** *unclassified*). Hence $P(db_1) = confidential$. For any policy G, the predicate $m(G, P)$ holds if db_1 does not leak confidential information (i.e. there is no non-reflexive outgoing edge from db_1).

Security invariants formalize security goals. A template contributes the formal semantics. A host mapping contains the scenario-specific knowledge. This makes the scenario-independent semantics available for formal reasoning by treating P and G as unknowns. Even reasoning with arbitrary security invariants is possible by additionally treating m as unknown.

With this modeling approach, the end user needs not to be bothered with the formalization of m, but only needs to specify G and P. In the course of this paper, we present a convenient method for specifying P.

Security Strategies and Monotonicity. In IT security, one distinguishes between two main classes of security strategies: *Access Control Strategies* (ACS) and *Information Flow Strategies* (IFS) [12, §6.1.4]. An IFS focuses on confiden-

*a predicate is a total, Boolean-valued function.

tiality and an ACS on integrity or controlled access. We require that m is in one of these classes*.

The two security strategies have one thing in common: they prohibit illegal actions. From an integrity and confidentiality point of view, prohibiting more never has a negative side effect. Removing edges from the policy cannot create new accesses and hence cannot introduce new access control violations. Similarly, for an IFS, by statically prohibiting flows in the network, no new direct information leaks nor new side channels can be created. In brief, prohibiting more does not harm security. From this, it follows that if a policy (V, E) fulfills its security invariant, for a stricter policy rule set $E' \subseteq E$, the policy (V, E') must also fulfill the security invariant. We call this property *monotonicity*.

Composition of Security Invariants. Usually, there is more than one security invariant for a given scenario. However, composition and modularity is often a non-trivial problem. For example, access control lists that are individually secure can introduce security breaches under composition [13]. Also, information flow security of individually secure processes, systems, and networks may be subverted by composition [19]. This is known as the *composition problem* [2].

With the formalization in this paper, composability and modularity are enabled by design. For a fixed policy G with k security invariants, let m_i be the security invariant template and P_i the host mapping, for $i \in \{1 \dots k\}$. The predicate $m_i(G, P_i)$ holds if and only if the security invariant i holds for the policy G. With this modularity, composition of all security invariants is straightforward[i]: all security invariants must be fulfilled. The monotonicity guarantees that having more security invariants provides greater or equal security.

$$m_1(G, P_1) \wedge \dots \wedge m_k(G, P_k)$$

3 Examples of Security Invariant Templates

In this section, we present three examples of security invariant templates. Our implementation currently features more than ten templates and grows. All can be inspected in the published theory files. Common networking scenarios such as subnets, non-interference invariants, or access control lists are available. With the following templates, a larger case study is presented in Section 6.

Simplified Bell LaPadula with Trust. A simplified version of the Bell LaPadula model is already outlined in Example 1. In this paragraph, we extend this model with a notion of trust by adding a Boolean flag *trust* to the host attributes. For a host v, let $P(v).sc$ denote v's security clearance and $P(v).trust$ if

*By limiting m to IFS or ACS, we emphasize that availability is not in the scope of this work. Availability requires reasoning on a lower abstraction level, for example, to incorporate network hardware failure. Availability invariants could be expressed similarly, but would require inverse monotonicity (see below).

v is trusted. A trusted host can receive information of any security clearance and may declassify it, i. e. distribute the information with its own security clearance. For example, a trusted host is allowed to receive any information and with the *unclassified* clearance, it is allowed to reveal it to anyone. The template is thus formalized as follows.

$$m\big((V,E),\,P\big) \;\equiv\; \forall (s,r) \in E.\; \begin{cases} \mathit{True} & \textbf{if}\;\; P(r).\mathit{trust} \\ P(s).\mathit{sc} \;\leq\; P(r).\mathit{sc} & \textbf{otherwise} \end{cases}$$

Domain Hierarchy. The domain hierarchy template mirrors hierarchical access control structures. It is best introduced by example. The tiny car company

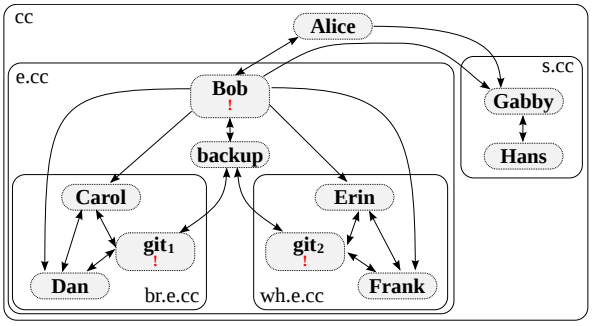

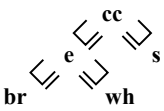

Fig. 2a (top): Organizational structure of **cc**.

Fig. 2b (left): Policy and host mapping of **cc**.

(cc) consists of the two sub-departments engineering (e) and sales (s). The engineering department itself consists of the brakes (br) and the wheels (wh) department. This tree-like organizational structure is illustrated in Fig. 2a. We denote a position by the fully qualified domain name, e. g., $wh.e.cc$ uniquely identifies the wheels department. Let '$\sqsubseteq$' denote the '*is below or at the same hierarchy level*' relation, e. g., $wh.e.cc \sqsubseteq wh.e.cc$, $wh.e.cc \sqsubseteq e.cc$, and $wh.e.cc \sqsubseteq cc$. However, $wh.e.cc \not\sqsubseteq br.e.cc$ and $br.e.cc \not\sqsubseteq wh.e.cc$. The '$\sqsubseteq$' relation denotes a partial order[ii]. The company's command structures are strictly hierarchical, i. e. commands are either exchanged in the same department or travel from higher departments to their sub-departments. Formally, the receiver's level $\sqsubseteq$ sender's level. For a host v, let $P(v).\mathit{level}$ map to the fully qualified domain name of v's department. For example in Fig. 2b, $P(\mathit{Bob}).\mathit{level} = e.cc$.

As in many real-world applications of a mathematical model, exceptions exist. Those are depicted by exclamation marks in Fig. 2b. For example, Bob as head of engineering is in a trusted position. This means he can operate as if he were in the position of Alice. This implies that he can communicate on par with Alice, which also implies that he might send commands to the sales department. We model such exceptions by assigning each host a trust level. This trust level specifies up to which position in the hierarchy this host may act. For example, Bob in $e.cc$ with a trust level of 1 can act as if he were in cc, which means he has the same command power as Alice. Let $P(v).\mathit{trust}$ map to v's trust level. We implement

a function $chop(level : DomainName,\ trust : \mathbb{N}) \Rightarrow DomainName$ which chops off *trust* sub-domains from a domain name, e.g., $chop(br.e.cc,\ 1) = e.cc$. With this, the security invariant template can be formalized as follows.

$$m\big((V,E),\ P\big)\ \equiv\ \forall (s,r) \in E.\ \ P(r).level\ \sqsubseteq\ chop\big(P(s).level,\ P(s).trust\big)$$

Security Gateway. Hosts may belong to a certain domain. Sometimes, a pattern where intra-domain communication between domain members must be approved by a central instance is required. As an example, let several virtual machines belong to the same domain and a secure hypervisor manage intra-domain communication. As another example, inter-device communication of slave devices in the same domain is controlled by a central master device. We call such a central instance 'security gateway' and present a template for this architecture. Four host roles are distinguished: A security gateway (*sgw*), a security gateway accessible from the outside (*sgwa*), a domain member (*memb*), and a default value that reflects 'none of these roles' (*default*). The following table implements the access control restrictions. The role of the sender (snd), role of the receiver (rcv), the result (rslt), and an explanation are given.

snd	rcv	rslt	explanation
sgw	*	✓	Can send to the world.
sgwa	*	✓	— "—
memb	*sgw*	✓	Can contact its security gateway.
memb	*sgwa*	✓	— "—
memb	*memb*	✗	Must not communicate directly. May communicate via $sgw(a)$.
memb	*default*	✓	No restrictions for direct access to outside world. Outgoing accesses are not within the invariant's scope.
default	*sgw*	✗	Not accessible from outside.
default	*sgwa*	✓	Accessible from outside.
default	*memb*	✗	Protected from outside world.
default	*default*	✓	No restrictions.

This template is minimalistic in that it only restricts accesses to members (from other members or the outside world), whereas accesses from members to the outside world are unrestricted. It can be implemented by a simple table lookup. In-host communication is allowed by adding $s \neq r$.

$$m\big((V,E),\ P\big)\ \equiv\ \forall (s,r) \in E,\ s \neq r.\ \ \texttt{table}\big(P(s),\ P(r)\big)$$

4 Generic Semantic Analysis of Security Invariants

Offending Flows. Since m is monotonic, if an IFS or ACS security invariant is violated, there must be some flows in G that are responsible for the violation. By removing them, the security invariant should be fulfilled (if possible). We

call a minimal set of such flows the *offending flows*. Minimality is expressed by requiring that every single flow in the offending flows bears responsibility for the security invariant's violation.

Definition 3 (Set of Offending Flows).

$$\begin{aligned} \mathit{set_offending_flows}\big(G,\, P\big) = \big\{ F \subseteq E \mid \neg\, m\big(G,\, P\big) \;\wedge\; m\big((V,\; E \setminus F),\, P\big) \;\wedge \\ \forall (s,r) \in F.\; \neg\, m\big((V,\; (E \setminus F) \cup \{(s,r)\})\,,\, P\big) \big\} \end{aligned}$$

Example 2. The definition does not require that the offending flows are uniquely defined. This is reflected in its type since it is a set of sets. For example, for $G = (\{v_1, v_2, v_3\}, \{(v_1, v_2), (v_2, v_3)\})$ and a security invariant that v_1 must not transitively access v_3, the invariant is violated: v_2 could forward requests. The set of offending flows is $\{\{(v_1, v_2)\}, \{(v_2, v_3)\}\}$. This ambiguity tells the end user that there are multiple options to fix a violated security invariant. The policy can be tightened by prohibiting one of the offending flows, e. g.,$\{(v_1, v_2)\}$.

If $m(G,\, P)$ holds, the set of offending flows is always empty[iii]. Also, for every element in the set of offending flows, it is guaranteed that prohibiting these flows leads to a fulfilled security invariant[iv]. It is not guaranteed that the set of offending flows is always non-empty for a violated security invariant. Depending on m, it may be possible that no set of flows satisfies Def. 3. However, Theorem 1 proves[v] an important insight: a violated invariant can always be repaired by tightening the policy if and only if the invariant holds for the deny-all policy.

Theorem 1 (No Edges Validity). *For m monotonic, arbitrary V, E, and P, let $G = (V,\, E)$ and $G_{\mathit{deny\text{-}all}} = (V, \emptyset)$. If $\neg\, m(G, P)$ then*

$$m\big(G_{\mathit{deny\text{-}all}},\, P\big) \longleftrightarrow \mathit{set_offending_flows}(G,\, P) \neq \emptyset$$

We demand that all security invariants fulfill $m\big(G_{\mathit{deny\text{-}all}},\, P\big)$. This means that violations are always fixable.

We call a host responsible for a security violation the *offending host*. Given one offending flow, the violation either happens at the sender's or the receiver's side. The following difference between ACS and IFS invariant can be observed. If m is an ACS, the host that initiated the request provokes the violation by violating an access control restriction. If m is an IFS, the information leak only occurs when the information reaches the unintended receiver. This distinction is essential as it renders the upcoming Def. 5 and 6 provable.

Definition 4 (Offending Hosts). *For $F \in \mathit{set_offending_flows}(G, P)$*

$$\mathit{offenders}(F) = \begin{cases} \{\, s \mid (s,r) \in F\} & \textbf{if } \mathit{ACS} \\ \{\, r \mid (s,r) \in F\} & \textbf{if } \mathit{IFS} \end{cases}$$

Secure Auto Completion of Host Mappings. Since P is a *total* function $\mathcal{V} \Rightarrow \Psi$, a host mapping for *every* element of $\mathcal{V}$ must be provided. However, an end user might only specify the *security-relevant* host attributes. Let $P_C \subseteq \mathcal{V} \times \Psi$ be a finite, possibly incomplete host attribute mapping specified by the end user. For some $\bot \in \Psi$, the total function P can be constructed by $P(v) \equiv$ (**if** $(v, \psi) \in P_C$ **then** ψ **else** $\bot$). Intuitively, if no host attribute is specified by the user, $\bot$ acts as a default attribute.

Given the user specified all security-relevant attributes, we observe that the default attribute can never solve an existing security violation. Therefore, we conclude that for a given security invariant m, a value $\bot$ can securely be used as a default attribute if it cannot mask potential security risks. In other words, a default attribute $\bot$ is secure w.r.t. the given information P if for all offenders v, replacing v's attribute* by $\bot$, denoted by $P_{v \mapsto \bot}$, has the same amount of security-relevant information as the original P.

Definition 5 (Secure Default Attribute). *A $\bot$ is a secure default attribute iff for a fixed m and for arbitrary G and P that cause a security violation, replacing the host attribute of any offenders by $\bot$ must guarantee that no security-relevant information is masked.*

$$\forall\, G\ P.\ \forall\, F \in \mathit{set_offending_flows}\big(G, P\big).\ \ \forall v \in \mathit{offenders}(F).\ \neg\, m\big(G, P_{v \mapsto \bot}\big)$$

Example 3. In the simple Bell LaPadula model, an IFS, let us assume information is leaked. The predicate 'information leaks' holds, no matter to which lower security clearance the information is leaked. In general, if there is an illegal flow, it is from a higher security clearance at the sender to a lower security clearance at the receiver. Replacing the security clearance of the receiver with the lowest security clearance, the information about the security violation is always preserved. Thus, *unclassified* is the secure default attribute[vi]. In summary, if all classified hosts are labeled correctly, treating the rest as unclassified prevents information leakage.

To elaborate on Def. 5, it can be restated as follows. It focuses on the available security-relevant information in the case of a security violation. The attribute of an offending host v bears no information, except for the fact that there is a violation. A secure default attribute $\bot$ cannot solve security violations. Hence $P(v)$ and $\bot$ are equal w.r.t. the security violation. Thus, P and $P_{v \mapsto \bot}$ must be equal w.r.t. the information about the security violation. Requiring this property for all policies, all possible security violations, all possible choices of offending flows, and all candidates of offending hosts, this definition justifies that $\bot$ never hides a security problem.

Example 4. Definition 5 can be specialized to the exemplary case in which a new host x is added to a policy G without updating the host mapping. Consulting an oracle, x's real host attribute is $P(x) = \psi$. In reality, the oracle is not available and x is mapped to $\bot$ because it is new and unknown. Let x be an attacker.

*$P_{v \mapsto \bot} \equiv (\lambda x.$ **if** $x = v$ **then** $\bot$ **else** $P(x))$, an updated P which returns $\bot$ for v

With the oracle's ψ-attribute, x causes a security violation. We demand that the security violation is exposed even without the knowledge from the oracle. Definition 5 satisfies this demand: if x mapped by the oracle to ψ causes a security violation, x mapped to $\bot$ does not mask the security violation.

A 'deny-all' default attribute is easily proven secure. Definition 5 reads the following for this case: if an offender v does something that violates $m(G, P)$, then removing all of v's rights ($P_{v \mapsto deny\text{-}all}$), a violation must persist. Hence, designing whitelisting security invariant templates with a restrictive default attribute is simple. However, to add to the ease-of-use, more permissive default attributes are often desirable since they reduce the manual configuration effort. In particular, if a security invariant only concerns a subset of a policy's hosts, no restrictions should be imposed on the rest of the policy. This is also possible with Def. 5, but may require a comparably difficult proof.

Example 5. In Example 1, no matter how many hosts are added to the policy, it is sufficient to only specify that db_1 is *confidential*. This confidentiality is guaranteed while no restrictions are put on hosts that do not interact with db_1.

Definition 6 (Default Attribute Uniqueness). *A default attribute $\bot$ is called unique iff it is secure (Def. 5) and there is no $\bot' \neq \bot$ s.t. $\bot'$ is secure.*

We demand that all security invariants fulfill Def. 6. This means that there is only one unique secure default attribute $\bot$.

Example 6. In the simple Bell LaPadula model, since the security clearances form a total order, the lowest security clearance is uniquely defined.

With the experience of proving Def. 5 and 6 for default attributes for 18 invariant templates, the connection between offending host and security strategy was discovered. During our early research, we realized that a Boolean variable, fixed for m, indicating the offending host was necessary to make Def. 5 and 6 provable. A classification of the different invariants revealed the important connection.

Default Attributes of Section 3's Templates. In the Bell LaPadula with Trust template, the default attribute is $(\mathit{unclassified}, \mathit{untrusted})$[vii]. In the Domain Hierarchy, it is[viii] a special value $\bot$ with a trust of zero and which is at the lowest point in the hierarchy, i.e. $\forall\, l.\ \bot \sqsubseteq l$. Finally, it is worth mentioning that the $\sqsubseteq$-relation forms a lattice[ix], which is a desirable structure for security classes [10]. In the Security Gateway, the default attribute is *default*[x].

All default attributes allow flows between each other. This greatly adds to the ease-of-use, since the scope of an invariant is limited only to the explicitly configured hosts. The unconcerned parts of a security policy are not negatively affected.

Unique and Efficient Offending Flows. All security invariant templates presented in Section 3 have a simple, common structure: a predicate is evaluated for all flows. Let $\Phi(\Psi, \Psi)$ be this predicate. Note that all invariants of this structure fulfill monotonicity[xi].

Since Def. 3 is defined over all subsets, the naive computational complexity of is in ***NP***. This section shows that – with knowledge about a concrete security invariant template m – it can be computed in linear time. For Φ-structured invariants, the offending flows are always uniquely defined and can be described intuitively[xii]*.

Theorem 2 (Φ Set of Offending Flows). *If m is Φ-structured $m(G, P) \equiv \forall (s, r) \in E.\ \Phi(P(s),\ P(r))$, then*

$$\mathit{set_offending_flows}\big(G,\ P\big) = \begin{cases} \{\{(s,r) \in E \mid \neg\ \Phi(P(s), P(r))\}\} & \textbf{if } \neg\ m(G,\ P) \\ \emptyset & \textbf{if } m(G,\ P) \end{cases}$$

Example 7. For the Bell LaPadula model, if no security violation exists the set of offending flows is $\emptyset$, else $\{\{(s, r) \in E \mid P(s) > P(r)\}\}$[xiv].

Policy Construction. A policy that fulfills all security invariants can be constructed by removing all offending flows from the allow-all policy $G_{all} = (V, V \times V)$. This approach is sound[xv] for arbitrary m and even complete[xvi] for Φ-structured security invariant templates.

Example 8. If completely contradictory security invariants are given, the resulting (maximum) policy is the deny-all policy $G_{deny\text{-}all} = (V, \emptyset)$.

5 Implementation

We built a tool called `topoS` with all the features presented in this paper. Its core reasoning logic consists of code generated by Isabelle/HOL. This guarantees the correctness of all results computed by `topoS`'s core [16].

Computational Complexity. `topoS` performs linear in the number of security invariants and quadratic in the number of hosts for Φ-structured invariants. For scenarios with less than 100 hosts, it responds interactively in less than 10 seconds. A benchmark of the automated policy construction, the most expensive algorithm, is presented in Fig. 3. For $|V|$ hosts, $|V|^2/4$ flows were created. With reasonable memory consumption, policies with up to 250k flows can be processed in less than half an hour. `topoS` contains a lot of machine generated code that is not optimized for performance but correctness. However, the overall theoretical and practical performance is sufficient for real-world usage. During our work with Airbus Group, we never encountered any performance issues.

*The same holds for templates with a structure similar to the Security Gateway[xiii]

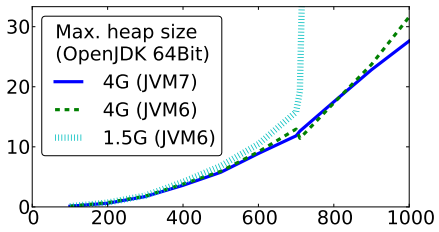

Fig. 3: Runtime of the policy construction algorithm for 100 Φ-structured invariants on an i7-2620M CPU (2.70GHz), Java Virtual Machine. X-axis: $|V|$, Y-axis: runtime in minutes.

Section 6 User Case Study: Statistics

Valid	Violations	Missing
16.0/15.8/1.7	1.0/3.2/4.4	5.0/5.6/2.1
14.0/14.0/1.4	1.0/1.6/1.9	7.0/7.4/1.0
12.0/10.6/5.7	4.0/6.6/4.9	11.0/11.2/6.2

median/arithmetic mean/std deviation

Table 1: Statistics on Section 6's user-designed policies. Number of valid, violating, and missing flows. User experiences (top to bottom): Expert, Intermediate, Novice. Five participants each.

6 Case Study: A Cabin Data Network

In this section, we present a slightly more complex scenario: a policy for a cabin data network for the general civil aviation. This example was chosen as security is very important in this domain and it provides a challenging interaction of different security invariants. It is a small imaginary toy example, developed in collaboration with Airbus Group. To make it self-contained and accessible to readers without aeronautical background knowledge, it does not obey aeronautical standards. However, the scenario is plausible, i. e. a real-world scenario may be similar. During our research, we also evaluated real world scenarios in this domain. With this experience, we try to present a small, simplified, self-contained, plausible toy scenario that, however, preserves many real world snares.

To estimate the scenario's complexity, we asked 15 network professionals to design its policy. On the one hand, as many use cases as possible should be fulfilled, on the other hand, no security violation must occur. Therefore, the task was to maximize the allowed flows without violating any security invariant. The results are illustrated in Table 1. Surprisingly, even expert network administrators made errors (both missing flows and security violations) when designing the policy.

A detailed scenario description, the host attribute mappings, and raw data are available in [11]. Using this reference, we also encourage the active reader to design the policy by oneself before it is revealed in Fig. 4b. The scenario is presented in the following compressed two paragraphs.

The network consists of the following hosts.

CC The Cabin Core Server, a server that controls essential aircraft features, such as air conditioning and the wireless and wired telecommunication of the crew.

C1, C2 Two mobile devices for the crew to help them organize, e. g., communicate, make announcements.

Wifi A wifi hotspot that allows passengers to access the Internet with their own devices. Explicitly listed as it might also be responsible for billing passenger's Internet access.

IFEsrv The In-Flight Entertainment server with movies, Internet access, etc. Master of the IFE displays.

IFE1, IFE2 Two In-Flight Entertainment displays, mounted at the back of passenger seats. They provide movies and Internet access. Thin clients, everything is streamed from the IFE server.

P1, P2 Two passenger-owned devices, e. g., laptops, smartphones.

Sat A satellite uplink to the Internet.

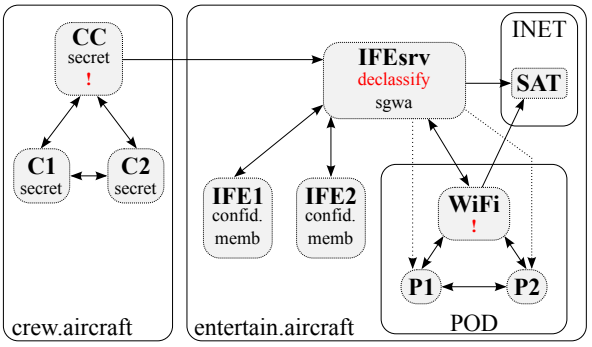

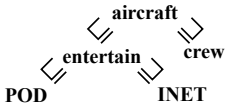

Fig. 4a (top): Security domains of the cabin data network.

Fig. 4b (left): Cabin network policy and hosts' attributes.

The following three security invariants are specified.

Security Invariant 1, Domain Hierarchy. Four different security domains exist in the aircraft, c.f. Fig. 4a. They separate the **crew** domain, the **entertain**ment domain, the passenger-owned devices (**POD**) domain and the Internet (**INET**) domain. In Fig. 4b, the host domain mapping is illustrated and trusted devices are marked with an exclamation mark.
The CC may send to the entertain domain, hence it is trusted. Possible use cases: Stewards coordinate food distribution; Announcement from the crew is send to the In-Flight Entertainment system (via CC) and distributed there to the IFE displays.
The Wifi is located in the **POD** domain to be reachable by PODs. It is trusted to send to the entertain domain. Possible use case: Passenger subscribes a film from the IFE server to her notebook or establishes connections to the Internet.
In the **INET** domain, the SAT is isolated to prevent accesses from the Internet into the aircraft.

Security Invariant 2, Security Gateway. The IFE displays are thin clients and strictly bound to their server. Peer to peer communication is prohibited. The Security Gateway model directly provides the respective access control restrictions.

Security Invariant 3, Bell LaPadula with Trust. Invariant 3 defines information flow restrictions by labeling confidential information sources. To protect the passenger's privacy when using the IFE displays, it is undesirable that the IFE displays communicate with anyone, except for the IFEsrv. Therefore, the IFE displays are

marked as confidential (confid.). The IFEsrv is considered a central trusted device. To enable passengers to surf the Internet on the IFE displays by forwarding the packets to the Internet or forward announcements from the crew, it must be allowed to declassify any information to the default (i. e. unclassified) security clearance. Finally, the crew communication is considered more critical than the convenience features, therefore, CC, C1, and C2 are considered secret. As the IFEsrv is trusted, it can receive and forward announcements from the crew.

This case study illustrates that this complex scenario can be divided into three security invariants that can be represented with the help of the previously presented templates. Figure 4b also reveals that very few host attributes must be manually specified; the automatically added secure default attributes are not shown. All security invariants are fulfilled.

The automated policy construction yields the following results. The solid edges unified with the dashed edges* result in the uniquely defined policy with the maximum number of allowed flows. The solid lines were given by the policy, the dashed lines were calculated from the invariants. These 'diffs' are computed and visualized automatically by `topoS`. They provide the end user with helpful feedback regarding '*what do my invariants require?*' vs. '*what does my policy specify?*'. This results in a feedback loop we used extensively during our research to refine the policy and the invariants. It provides a 'feeling' for the invariants.

The main evaluation of this work are the formal correctness proofs. However, we also presented `topoS` to the 15 users and asked them to personally judge `topoS`'s utility. It was considered downright helpful and a majority would want to use it for similar tasks. The graphical feedback was also much appreciated.

7 Related Work

In a field study with 38 participants, Hamed and Al-Shaer discovered that "even expert administrators can make serious mistakes when configuring the network security policy" [17]. Our user feedback session extends this finding as we discovered that even expert administrators can make serious mistakes when *designing* the network security policy.

In their inspiring work, Guttman and Herzog [15] describe a formal modeling approach for network security management. They suggest algorithms to verify whether configurations of firewalls and IPsec gateways fulfill certain security goals. These comparatively low-level security goals may state that a certain packet's path only passes certain areas or that packets between two networked hosts are protected by IPsec's ESP confidentiality header. This allows reasoning on a lower abstraction level at the cost of higher manual specification and configuration effort. Header space analysis [18] allows checking static network invariants such as no-forwarding-loops or traffic-isolation on the forwarding and middleboxes plane. It provides a common, protocol-agnostic framework and algebra on the packet header bits.

*unified with all reflexive edges, i. e. in-host communication

Firmato [1] was designed to ease management of firewalls. A firewall-independent entity relationship model is used to specify the security policy. With the help of a model compiler, such a model can be translated to firewall configurations. Ethane [6] is a link layer security architecture which evolved to the network operating system NOX [14]. They implement high-level security policies and propose a secure binding from host names to network addresses. In the long term, we consider `topoS` a valuable add-on on top of such systems for policy verification. For example, it could warn the administrator that a recent network policy change violates a security invariant, maybe defined years ago.

Expressive policy specification languages, such as Ponder [9], were proposed. Positive authorization policies (only a small aspect of Ponder) are roughly comparable to our policy graph. The authors note that e.g., negative authorization policies (deny-rules) can create conflicts. Policy constraints can be checked at compile time. In [4], a policy specification language (SPSL) with allow and deny policy rules is presented. With this, a conflict-free policy specification is constructed. Conflict-free Boolean formulas of this policy description and the policy implementation in the security mechanisms (router ACL entries) are checked for equality using a SAT solver. One unique feature covered is hidden service access paths, e.g., http might be prohibited in zone1 but zone1 can ssh to zone2 where http is allowed. [8] focuses on policies in dynamic systems and their analysis. These papers require specification of the verification goals and security goals and can thus benefit from our contributions.

This work's modeling concept is very similar to the Attribute Based Access Control (ABAC) model [21], though the underlying formal objects differ. ABAC distinguishes subjects, resources, and environments. Attributes may be assigned to each of these entities, similar to our host mappings. The ABAC policy model consists of positive rules which grant access based on the assigned attributes, comparably to security invariant templates. Therefore, our insights and contributions are also applicable to the ABAC model.

8 Conclusion

After more than 50k changed lines of formal theory, our simple, yet powerful, model landscape emerged. Representing policies as graphs makes them visualizable. Describing security invariants as total Boolean-valued functions is both expressive and accessible to formal analysis. Representing host mappings as partial configurations is end-user-friendly, transforming them to total functions makes them handy for the design of templates. With this simple model, we discovered important universal insights on security invariants. In particular, the transformation of host mappings and a simple sanity check which guarantees that security policy violations can always be resolved. This provides deep insights about how to express verification goals. The full formalization in the Isabelle/HOL theorem prover provides high confidence in the correctness.

Acknowledgments & Availability. We thank all the participants of our feedback session and the anonymous reviewers very much. A special thanks goes to our colleague Lothar Braun for his valuable feedback. Lars Hupel helped finalizing the paper. This work has been supported by the German Federal Ministry of Education and Research (BMBF) under support code 16BY1209F, project ANSII, and 16BP12304, EUREKA project SASER, and by the European Commission under the FP7 project EINS, grant number 288021.

Our Isabelle/HOL theory files and `topoS`'s Scala source code are available at

`https://github.com/diekmann/topoS`

References

1. Bartal, Y., Mayer, A., Nissim, K., Wool, A.: Firmato: A novel firewall management toolkit. In: IEEE Symposium on Security and Privacy. pp. 17–31. IEEE (1999)
2. Bell, D.: Looking back at the Bell-La Padula model. In: Proceedings of the 21st Annual Computer Security Applications Conference. pp. 337–351 (Dec 2005)
3. Bell, D., LaPadula, L.: Secure computer systems: A mathematical model. MTR-2547 Vol. II, The MITRE Corporation, Bedford, MA (May 1973)
4. Bera, P., Ghosh, S., Dasgupta, P.: Policy based security analysis in enterprise networks: A formal approach. IEEE Transactions on Network and Service Management 7(4), 231–243 (2010)
5. Bishop, M.: What is computer security? IEEE Security & Privacy 1 (Feb 2003)
6. Casado, M., Freedman, M.J., Pettit, J., Luo, J., McKeown, N., Shenker, S.: Ethane: taking control of the enterprise. In: Proceedings of the 2007 conference on Applications, technologies, architectures, and protocols for computer communications. pp. 1–12. SIGCOMM '07, ACM, New York, NY, USA (2007)
7. Common Criteria: Security assurance components. Common Criteria for Information Technology Security Evaluation CCMB-2012-09-003 (Sep 2012), `http://www.commoncriteriaportal.org/files/ccfiles/CCPART3V3.1R4.pdf`
8. Craven, R., Lobo, J., Ma, J., Russo, A., Lupu, E., Bandara, A.: Expressive policy analysis with enhanced system dynamicity. In: Proceedings of the 4th International Symposium on Information, Computer, and Communications Security. pp. 239–250. ASIACCS '09, ACM, New York, NY, USA (2009)
9. Damianou, N., Dulay, N., Lupu, E., Sloman, M.: The ponder policy specification language. In: Sloman, M., Lupu, E., Lobo, J. (eds.) Policies for Distributed Systems and Networks, Lecture Notes in Computer Science, vol. 1995, pp. 18–38. Springer Berlin Heidelberg (2001)
10. Denning, D.: A lattice model of secure information flow. Communications of the ACM 19(5), 236–243 (May 1976)
11. Diekmann, C., Hanka, O., Posselt, S.A., Schlatt, M.: Imaginary aircraft cabin data network (toy example) (Jul 2013), `http://www.net.in.tum.de/pub/diekmann/cabin_data_network.pdf`
12. Eckert, C.: IT-Sicherheit: Konzepte-Verfahren-Protokolle. Oldenbourg Verlag, 8th edn. (2013), ISBN 3486721380
13. Gong, L., Qian, X.: The complexity and composability of secure interoperation. In: Proceedings of 1994 IEEE Computer Society Symposium on Research in Security and Privacy. pp. 190–200 (1994)

14. Gude, N., Koponen, T., Pettit, J., Pfaff, B., Casado, M., McKeown, N., Shenker, S.: NOX: towards an operating system for networks. SIGCOMM Comput. Commun. Rev. 38(3), 105–110 (Jul 2008)
15. Guttman, J.D., Herzog, A.L.: Rigorous automated network security management. International Journal of Information Security 4, 29–48 (2005)
16. Haftmann, F., Nipkow, T.: Code generation via higher-order rewrite systems. In: Blume, M., Kobayashi, N., Vidal, G. (eds.) Functional and Logic Programming, Lecture Notes in Computer Science, vol. 6009, pp. 103–117. Springer Berlin Heidelberg (2010)
17. Hamed, H., Al-Shaer, E.: Taxonomy of conflicts in network security policies. In IEEE Communications Magazine 44(3), 134 – 141 (Mar 2006)
18. Kazemian, P., Varghese, G., McKeown, N.: Header space analysis: static checking for networks. In: Networked Systems Design and Implementation. NSDI'12, USENIX Association, Berkeley, CA, USA (2012)
19. McCullough, D.: A hookup theorem for multilevel security. Software Engineering, IEEE Transactions on 16(6), 563 –568 (Jun 1990)
20. Nipkow, T., Paulson, L.C., Wenzel, M.: Isabelle/HOL: A Proof Assistant for Higher-Order Logic, Lecture Notes in Computer Science, vol. 2283. Springer (2002, last updated 2013), ISBN 978-3-540-43376-7
21. Yuan, E., Tong, J.: Attributed based access control (ABAC) for web services. In: IEEE International Conference on Web Services (2005)

Definitions, Lemmata, and Theorems

[i]all-security-requirements-fulfilled [ii]instantiation domainNameDept :: order [iii]validmodel-imp-no-offending [iv]remove-offending-flows-imp-model-valid [v]valid-empty-edges-iff-exists-offending-flows [vi]interpretation BLPbasic: NetworkModel [vii]interpretation BLPtrusted: NetworkModel [viii]interpretation DomainHierarchyNG: NetworkModel [ix]instantiation domainName :: lattice [x]interpretation SecurityGatewayExtended-simplified: NetworkModel [xi]monotonicity-eval-model-mono [xii]ENF-offending-set [xiii]ENFnr-offending-set [xiv]BLP-offending-set [xv]generate-valid-topology-sound [xvi]generate-valid-topology-max-topo

Appendix

In this appendix, we provide the host attribute mappings of Section 6's case study.

Domain Hierarchy

CC	$\mapsto$	(level : *crew.aircraft*, trust : 1)
C1	$\mapsto$	(level : *crew.aircraft*, trust : 0)
C2	$\mapsto$	(level : *crew.aircraft*, trust : 0)
IFEsrv	$\mapsto$	(level : *entertain.aircraft*, trust : 0)
IFE1	$\mapsto$	(level : *entertain.aircraft*, trust : 0)
IFE2	$\mapsto$	(level : *entertain.aircraft*, trust : 0)
SAT	$\mapsto$	(level : *INET.entertain.aircraft*, trust : 0)
Wifi	$\mapsto$	(level : *POD.entertain.aircraft*, trust : 1)
P1	$\mapsto$	(level : *POD.entertain.aircraft*, trust : 0)
P2	$\mapsto$	(level : *POD.entertain.aircraft*, trust : 0)

Security Gateway

IFEsrv	$\mapsto$	*sgwa*
IFE1	$\mapsto$	*memb*
IFE2	$\mapsto$	*memb*

Simplified Bell LaPadula with Trust

CC	$\mapsto$	(sc : *secret*, trust : *False*)
C1	$\mapsto$	(sc : *secret*, trust : *False*)
C2	$\mapsto$	(sc : *secret* trust : *False*)
IFE1	$\mapsto$	(sc : *confidential*, trust : *False*)
IFE2	$\mapsto$	(sc : *confidential*, trust : *False*)
IFEsrv	$\mapsto$	(sc : *unclassified*, trust : *True*)